arXiv:1410.5078v1 [cs.AI] 19 Oct 2014

Learning Vague Concepts for the Semantic Web

Paolo Pareti and Ewan Klein

School of Informatics, University of Edinburgh
10 Crichton Street, Edinburgh EH8 9AB, UK
P.Pareti@sms.ed.ac.uk, ewan@inf.ed.ac.uk

Abstract. Ontologies can be a powerful tool for structuring knowledge, and they are currently the subject of extensive research. Updating the contents of an ontology or improving its interoperability with other ontologies is an important but difficult process. In this paper, we focus on the presence of vague concepts, which are pervasive in natural language, within the framework of formal ontologies. We will adopt a framework in which vagueness is captured via numerical restrictions that can be automatically adjusted. Since updating vague concepts, either through ontology alignment or ontology evolution, can lead to inconsistent sets of axioms, we define and implement a method to detecting and repairing such inconsistencies in a local fashion.

1 Introduction

Historically, there has been a close relationship between ontologies on the one hand, and glossaries, taxonomies and thesauri on the other hand: although formal ontologies are expressed in a well-defined formal language, many of the intuitions and terms used in ontologies are derived from their natural language counterparts. Nevertheless, there is an obvious mismatch between formal ontologies and natural language expressions: vagueness is pervasive in natural language, but is typically avoided or ignored in ontologies.

Following standard usage, we will say that a concept is **vague** when it admits borderline cases — that is, cases where we are unable to say whether the concept holds or fails to hold.[1] The standard example involves the adjective *tall*. There are some people that we regard as definitely tall, and others we regard as definitely short; but people of average height are neither tall nor short. Notice that the source of indeterminacy here is not lack of world knowledge: we can know that John is, say, 1.80 metres in height, and still be undecided whether he counts as tall or not.

Rather than trying to capture vague expressions directly (for example, by means of fuzzy logic), we will view vagueness as a property that characterizes the definition of certain concepts over a *sequence* of ontologies deployed by an agent. While 'ordinary' concepts are treated as having a fixed meaning, shared by all users of the ontology, we propose instead that the meaning of a vague

[1] For recent overviews of the very extensive literature on vagueness, see [31,32]

concept is *unstable*, in the sense that the threshold on the scale of height which distinguishes between being tall and not tall is inherently defeasible.

Is there any reason why we should care about ontologies being able to express vagueness? As an example, consider FOAF [6], which is one of the most widely used ontologies on the web. One of FOAF's core predicates is `based_near`. It is instructive to read the commentary on this property:

> The `based_near` relationship relates two "spatial things" (anything that can be somewhere), the latter typically described using the geo:lat / geo:long geo-positioning vocabulary...
>
> We do not say much about what 'near' means in this context; it is a 'rough and ready' concept. For a more precise treatment, see GeoOnion vocab design discussions, which are aiming to produce a more sophisticated vocabulary for such purposes.

The concept is 'rough and ready' in a number of senses: it is undeniably useful; it is vague in that there are obviously borderline cases; and it is also highly context-dependent. This latter issue is addressed to a certain extent by the GeoOnion document [34] which is referenced above, but there is no systematic attempt to get to grips with vagueness.

We have chosen to implement our approach in OWL 2 [33], since we are interested in targetting semantic applications on the web, and OWL 2 is sufficiently expressive for our purposes while offering efficient reasoners such as Pellet [29] and HermiT [28]. Since the relevant aspects of OWL 2 can also be expressed more compactly in Description Logic (DL) [3], we will use the latter as our main vehicle for representing ontologies.

In this paper, we will start off (§1) by considering how to accommodate vague concepts into a framework such as Description Logic, and we will also situate the discussion within the wider perspective of ontology evolution and ontology alignment. Then §2 presents the architecture of the VAGO system, which treats vague concepts as defeasible; that is, able to be updated when new information is acquired. §3 describes and discusses a number of experiments in which the implemented VAGO system runs with both artificial and real-world data. Finally, §4 provides a conclusion.

2 Representing and Updating Vague Concepts

2.1 Gradable Adjectives and Measures

Adjectives such as *tall*, *expensive* and *near* are often called **gradable**, in that they can be combined with degree modifiers such as *very* and have comparative forms (*taller*, *more expensive*). As a starting point for integrating such predicates into Description Logic, consider the following degree-based semantic representation of the predicate *expensive*

$$\textit{expensive} \equiv \lambda x.\exists d[\mathbf{C}(d) \wedge \mathbf{expensive}(x) \preceq d] \tag{1}$$

Here, "**expensive** represents a measure function that takes an entity and returns its cost, a degree on the scale associated with the adjective" [21], p.349. The predicate **C** is a contextually-given restriction which determines the threshold for things that are definitely expensive. Thus, an object will be expensive if its cost is greater than the threshold d. The relation **expensive** resembles a datatype property in OWL, associating an individual with a data value. In DL, we could introduce a concept $\mathsf{Expensive}$ and constrain its interpretation with a datatype restriction of the following kind [27], where X is a variable whose role we will discuss shortly:

$$\mathsf{Expensive} \sqsubseteq \exists \mathsf{hasMeasure}.(\geq, X) \tag{2}$$

In [2], an expression such as $(\geq, 200)$ is called a predicate name, and corresponds to an abstraction over a first-order formula, e.g., $\lambda x.x \leq 200$.

To gain a more general approach, we will adopt the approach to adjectives proposed in [1] and make the predicate's scale (in this case, the cost) explicit:

$$\mathsf{Expensive} \equiv \exists \mathsf{hasProperty}.(\mathsf{Cost} \sqcap \exists \mathsf{hasMeasure}.(\geq, X)) \tag{3}$$

However, we are left with a problem, since we still need some method of to provide a concrete value in place of X in particular contexts of use. We will regard X as a metavariable, and to signal its special function, we will call it an **adaptor**. Any DL axiom that contains one or more adaptors is an **axiom template**. Suppose $\phi[X_1, \ldots X_n]$ is an axiom template, $\mathcal{X} = \{X_1, \ldots X_n\}$ is the set of adaptors in ϕ and $T^{\mathcal{D}}$ is a set of datatype identifiers. A **assignment** $\theta : \mathcal{X} \mapsto T^{\mathcal{D}}$ is a function which binds a datatype identifier to an adaptor. We write $\phi[X_1, \ldots X_n]\theta$ for the result of applying the assignment θ to $\phi[X_1, \ldots X_n]$. For example, if $\phi[X]$ is the axiom template in (2), then $\phi[X]\{X \leftarrow 200\}$ is $\mathsf{Expensive} \sqsubseteq \exists \mathsf{hasMeasure}.(\geq, 200)$.

2.2 Delineations

Gradable adjectives typically come in pairs of opposite polarity; for example, the negative polarity opposite of *expensive* is *cheap*, which we can define as follows:

$$\mathsf{Cheap} \equiv \exists \mathsf{hasProperty}.(\mathsf{Cost} \sqcap \exists \mathsf{hasMeasure}.(\leq, X')) \tag{4}$$

Should the value of X' in (4) — the upper bound of definitely cheap — be the same as the value of X in (3) — the lower bound of definitely expensive? On a partial semantics for vague adjectives, the two will be different. That is, there be values between the two where things are not clearly expensive or cheap. One problem with partial valuation is that if $C(x)$ is neither true nor false for some x then plausible treatments of logical connectives would give the same undefined value to $C(x) \wedge \neg C(x)$ and $C(x) \vee \neg C(x)$. However, many logicians would prefer these propositions to retain their classical values of true and false respectively. To address this problem, Fine [11] and Kamp [20] proposed the use of **supervaluations**, namely valuations which make a partial valuation more

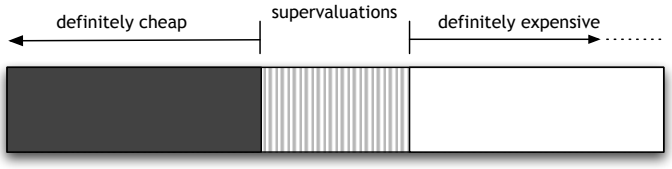

Fig. 1. Delineations of a Vague Concept

precise by extending them to a total valuation of the standard kind. In place of formal details, let's just consider a diagram: That is, each supervaluation can be thought of as way of finding some value v in the 'grey area' which is both an upper bound threshold for *cheap* and a lower bound threshold for *expensive*.

We adopt a model of vagueness in which vague concepts *do* receive total interpretations, but these are to be regarded as similar to supervaluations, in the sense that there may be multiple admissible delineations for the concept. The particular delineation that is adopted by an agent in a specific context can be regarded as the outcome of learning, or of negotiation with other agents. Consequently, on this approach, vague concepts differ from crisp concepts by only virtue of their instability: a vague concept is one where the threshold is always open to negotiation or revision. As we have already indicated, the defeasibility of threshold is not completely open, but is rather restricted to some borderline area; however, we will not attempt to formally capture this restriction here.[2]

2.3 Ontology Evolution and Ontology Alignment

As part of a learning process, an agent should be prepared to update the value of adaptors occurring in concepts in its ontology. New values can be learned as a result of interaction with other agents—corresponding to **ontology alignment** [10,8], or as a result of updating its beliefs about the world—corresponding to **ontology evolution** [12]. We will examine two important issues. The first concerns the question of how to automatically update vague concepts in an ontology as a result of learning new information [35]. The second, closely related issue, is how to ensure that an ontology is still consistent after some of its axioms have been modified [22].

Let's assume we have two ontologies $O_1 = (S, A_1)$ and $O_2 = (S, A_2)$ which share a signature S but have differing axiom sets A_1 and A_2. Axioms will consist of concept inclusions $C \sqsubseteq D$, concept assertions $C(a)$ and role assertions $r(a, b)$. Suppose A_1 contains the following axioms:

$$\mathsf{LegalAdult}(\mathsf{Jo}) \tag{5}$$

$$\mathsf{LegalAdult} \equiv \mathsf{Person} \sqcap \exists \mathsf{hasAge}.(\geq, 18) \tag{6}$$

[2] For more discussion, see [24,7].

We now want to update O_1 with the axiom set A_2, which happens to contain the following:

$$\mathsf{hasAge}(\mathsf{Jo}, 17) \tag{7}$$

How do we deal with the ensuing inconsistency? The core of our proposal involves identifying the numerical restriction in (1) as the value of an adaptor parameter, and therefore open to revision. That is, like other approaches to Ontology Repair, we need to identify and modify axioms that are responsible for causing an inconsistency. However, we localize the problem to one particular component of axioms that provide definitions for vague concepts. In this example, the inconsistency could be solved by changing the value 18 used in the definition of a LegalAdult to the value of 17.

3 System Architecture

3.1 Overview

As described earlier, vagueness is captured by the 'semantic instability' of certain concepts. This can be seen as an extreme kind of defeasibility: the threshold of a vague concept has a propensity to shift as new information is received. We have developed a a computational framework called VAGO in which learning leads to a change in the extension of vague concepts via the updating of adaptors. This is the only kind of ontology update that we will be considering here.

The first input to VAGO is the ontology $O = (S, A)$ that will potentially be updated. We call this the **original ontology** to distinguish it from the **updated ontology** which is the eventual output of the system. The second input is a set A_t of axioms, here called **training axioms**, that will be used to trigger an update of the original ontology. A_t is assumed to be part or all of the axiom set of some other ontology $O' = (S, A')$, and uses the same signature S as O.

A schematic representation of VAGO's architecture is shown in Fig 2.

Given the original ontology and the training axioms as inputs, the framework will output an updated version of the ontology. The whole process of computing this output can be divided into three main phases: validation, learning and update.

The goal of the validation phase is to extract diagnostic feedback from the training axioms. This feedback should provide information about the adaptors used in the original ontology. In particular, it should state whether an adaptor is responsible for an inconsistency and if so, propose an alternative value for the adaptor to remove the inconsistency.

The purpose of second phase, namely learning, is to determine how the values of adaptors should be updated, given the diagnostic feedback extracted during validation. These updates will in turn form the input of the third phase, which controls in detail how the original ontology should be modified to yield the updated ontology.

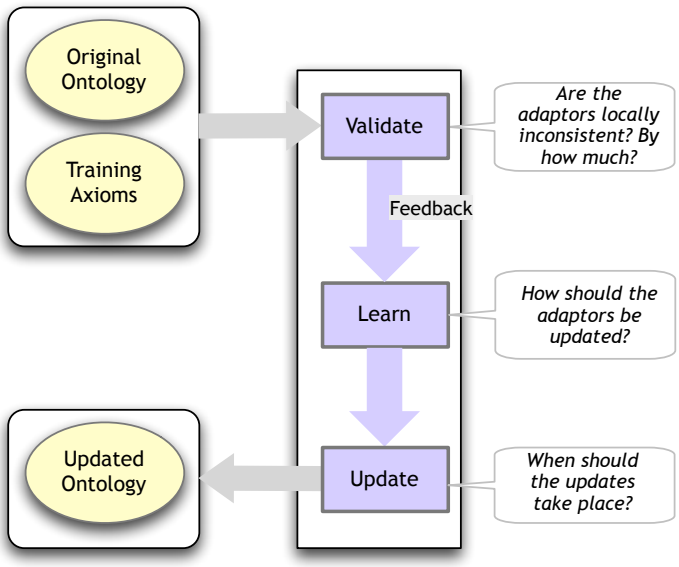

Fig. 2. Diagram of System Architecture

3.2 Validation Phase

Validation is the first phase of VAGO and examines the training axioms to produce a number of diagnostic **feedback objects** (FOs for short). These will contain a compact representation of all the information required for the subsequent learning phase. Let's suppose that in the original ontology, it is asserted that only people over the age of 18 are legal adults, where 18 is the value of adaptor X. In DL this assertion could be represented by the following axiom template:

$$\mathsf{Adult} \equiv \mathsf{Person} \sqcap \exists \mathsf{hasAge}.(\geq, X) \tag{8}$$

with the adaptor X instantiated to the value of 18. Table 1

Table 1. Example of feedback objects for adaptor X

	FO-1	FO-2
Adaptor identifier	X	X
Was the adaptor correct?	False	True
Current value of the adaptor	18	18
Required value for the adaptor	16	26

shows the feedback that the validation phase would output after examining the following training axioms:

$$\mathsf{LegalAdult}(\mathsf{John}) \tag{9}$$
$$\mathsf{hasAge}(\mathsf{John}, 16) \tag{10}$$
$$\mathsf{LegalAdult}(\mathsf{Jane}) \tag{11}$$
$$\mathsf{hasAge}(\mathsf{Jane}, 26) \tag{12}$$

In this example, training axioms (9) and (10) are incompatible with the definition of LegalAdult in the original ontology, generating an inconsistency. More specifically, under the assignment $X \leftarrow 18$, the set consisting of (8) and axioms (9), (10) is inconsistent, and in such a case we shall say that $X \leftarrow 18$ (or more briefly, just X) is **locally inconsistent**. However, this inconsistency can be removed by assigning a new value to X; more specifically, no inconsistency would arise for X holding a value less than or equal to 16. The optimal new value is one that removes the inconsistency with the least modification to the current value of X, which in this case is 16.

How is it possible to automatically determine, among all the possible values of X, the value v that will remove the inconsistency while differing least from the current value of X? Fortunately, if v exists, it is straightforwardly recoverable from training axioms. It is therefore sufficient to consider only a small set of possible candidate values for the adaptor. Given a set of inconsistent axioms (possibly minimal) and an adaptor X, algorithm 1.1 shows how to extract this candidate set.

Given the set V of values computed by Algorithm 1.1 for a set A of inconsistent axioms and an adaptor X, if there is an assignment $X \leftarrow v$ which will remove the inconsistency from axioms A, then v is either included in V, or it is the immediate predecessor or successor of a value in V.[3] For each value $v \in V$, it is necessary to consider also the first successor and first predecessor to deal with both strict and non-strict inequalities in numerical constraints.

3.3 Learning Phase

The learning phase is responsible for computing the updates that the original ontology should adopt, given the feedback objects extracted from the training axioms. The feedback objects are intended to provide the evidence necessary to justify a change in the adaptors.

If an adaptor assignment was found to be locally inconsistent, then it reasonable to assume that there is evidence to support its change. More specifically, given feedback to the effect that assignment $X \leftarrow v_0$ is locally inconsistent whereas $X \leftarrow v_1$ is not, then there is evidence to support a new assigment $X \leftarrow v_2$, where $v_2 = (v_1 - v_0)$.

[3] We define b to be the **immediate successor** of a if $a < b$ and there is no other value c such that $a < c < b$; the immediate prececessor is defined in a parallel manner.

Algorithm 1.1. Compute all candidate values for set A of inconsistent axioms and adaptor X

```
computeAlternativeValues(parameters: inconsistent_axioms,
    adaptor)
  values ← empty set
  data_relations ← set of relations in inconsistent_axioms
      restricted by the adaptor on value of their target
  cardinality_relations ← set of relations in
      inconsistent_axioms restricted by the adaptor in their
      cardinality
  all_individuals ← set of individuals in
      inconsistent_axioms
  foreach individual in all_individuals do
    foreach r in data_relations do
      data_values ← set of all values that individual is
          related to by relation r
      values ← values + all data_values
    end
    foreach r in cardinality_relations do
      cardinality ← number of relations r that the
          individual has
      values ← values + cardinality
    end
  end
  return (values)
```

The validation phase can discover different pieces of evidence supporting different values $[v_1, v_2, ..., v_n]$ for the same adaptor X. We will assume that the update to be computed should be the arithmetic mean of all these possible values. If the information contained in the training axioms is subject to noise, it will be desirable to reduce the importance of new values that are far from the mean $\bar{v}$. For this purpose, we use a sigmoid function $l : \mathbb{R} \mapsto [0, 1]$ to reduce exponentially the importance of a candidate value v the further it is from the mean $\bar{v}$. Values far from the mean will be scaled by a factor close to 0 (e.g., $l(\infty) = 0$) while those close to the mean will be scaled by a factor close to 1 (e.g., $l(0) = 1$). Given the distance x from the mean and the standard deviation δ over the set V of candidate values, a plausible definition for the function l might be the following (where q and b are free parameters):

$$l(x) = \frac{1}{\left(1 + (\delta/q)e^{-b(x-2\delta)}\right)^{q/\delta}}$$

After these additional considerations, the update u can be computed analytically using the following formula:

$$u = \frac{1}{n}\left(\sum_{i=1}^{n} v_i \, l(|v_i - \bar{v}|)\right)$$

3.4 Implementation

We have built a Java implementation of VAGO using the OWL API version 3.2.3 [18] to manipulate OWL 2 ontologies.[4] Adaptors are identified and modified by accessing the XML/RDF serialization of the ontology. The operations on the XML data make use of the JDOM API version 1.1.1 [19]. Additional information has to be added to the ontology in order to represent adaptors. To preserve the functionality of the OWL ontologies, any additional information required by VAGO can be encoded as axiom annotations, or as attributes in the RDF/XML serialization of the ontology. As a result, this information will be transparent to any reasoner and it will not change the standard semantics of the ontology.

Adaptors will be identified in OWL ontologies using special labels. More specifically, if a value v used in axiom A is labeled with an unique identifier associated with adaptor X, then it is possible to say that X is currently holding the value v and that the axiom A is *dependent* on the adaptor X. When the value of adaptor X is updated to a new value z, then the value v in axiom A will be replaced with the new value z. If multiple axioms of an ontology are dependent on adaptor X, then all their internal values associated with X will change accordingly.

4 Experiments

4.1 Experiments using Artificial Data

The first evaluation of VAGO uses an ontology describing persons and simple relations between them. The most important definitions in this ontology are the following:

- Person: a general class representing a human being;
- Minor $\equiv$ Person $\sqcap$ $\exists$hasAge.$(<, X_1)$: a class representing a young person (defined as a person under the age of X_1);
- LegalAdult $\equiv$ Person $\sqcap$ $\exists$hasAge.$(\geq, X_1)$: a class representing an adult (defined as a person of age X_1 or older);
- BusyParent $\equiv$ Person $\sqcap$ $\geq X_2$parentOf.Minor: a class representing the vague concept of a busy parent (defined as a person with at least X_2 young children);

[4] The code for the implementation is available at `http://bitbucket.org/ewan/vago/src`.

– RelaxedParent ≡ Person ⊓ ∃parentOf.Person ⊓ ¬BusyParent: a class representing the vague concept of a relaxed parent (defined as a parent that is not busy);
– hasAge: a functional data property with domain Person and range integer values;
– parentOf: an object relation between two Persons, a parent and a child.

The training axioms used in each iteration are produced automatically by generating instances of the above mentioned classes, and the relations between them. Since the data is produced automatically, it is possible to know the exact value that the adaptors should have, namely the value used while producing the data.

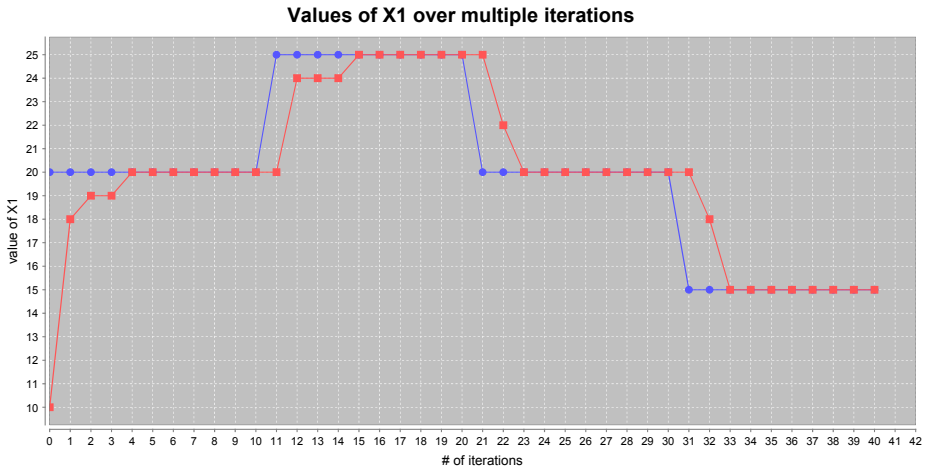

Fig. 3. Plot of the values of adaptor X_1 across 40 iterations of the system. The red squares indicate the value computed by the system, the blue circles indicate the correct value for that iteration.

In the first scenario, the system tries to adjust the adaptor X_1 that defines the threshold between the concepts Minor and LegalAdult by examining a number of individuals described in the training axioms. Forty sets of training axioms are generated, each one containing information about thirty persons. Whether these individuals are labelled as Minor or LegalAdult depends on their age (randomly determined by the generation algorithm) and on the correct value that X_1 should have in that iteration. The correct value for this adaptor changes every 10 iterations to test whether the system can adapt to changes in the environment. The results of this simulation are shown in Fig. 3. It can be seen that under these conditions the adaptor X_1 quickly converges to a consistent value.

The second scenario is concerned with the evolution of the adaptor X_2, which restricts the cardinality of a relation. In each of the 40 iterations of the system,

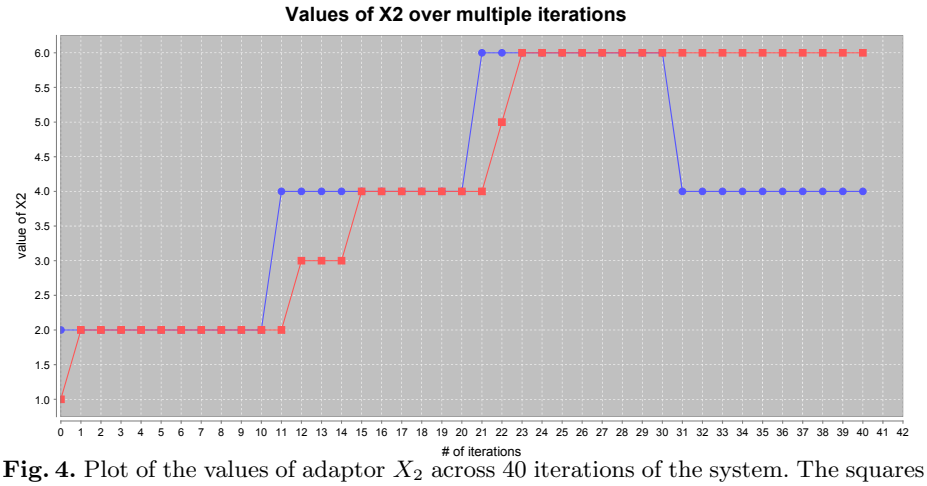

Fig. 4. Plot of the values of adaptor X_2 across 40 iterations of the system. The squares indicate the value computed by the system, the circles indicate the target value for that iteration.

a set of training axioms is generated so that the value of X_2 in that iteration is varied randomly. Each set of training axioms contains information about ten instances of the BusyParent class, such that they are parentOf at least X_2 children. It also contains ten instances of the class RelaxedParent, such that they are parentOf less than X_2 children.

Fig. 4 illustrates the result of the simulation in this second scenario, and shows that an adaptor restricting a cardinality (in this case X_2) can converge to a consistent value as long as the convergence involves an increase in the value. If, instead, the convergence *reduces* the value of the adaptor (as in the last ten iterations of this simulation), the adaptor will not change. For this reason, X_2 is not able to adapt to the last change in the simulation, namely reducing its value from 6 to 4. The inability to reduce the value of X_2, which might seem undesirable from a practical point of view, is a logical consequence of the Open World Assumption [9] made by the reasoner used.

4.2 Experiments with Web Data

This second evaluation deals with a hypothetical scenario where the user of VAGO is a company that owns bookshops in cities around the world. This company is interested in developing an automatic system to discover cities where it might be profitable to open a new bookshop by providing a sufficient definition for the concept 'profitable place'. We will here make the simplified assumption that a city centre will be classified as profitable place to open a business if there are few competitor bookshops nearby. In the centre of a city, potential customers of books are assumed to reach a bookshop on foot. Therefore in this context the concept 'near' should be interpreted as within walking distance. However, even after this clarification, two concepts remain underspecified. How many bookshops

should there be in a city centre to be able to say that they are "too many"? And how many meters away should an object be to count as "near"?

These vague concepts are defined in an OWL ontology using two adaptors. The first one, C, determines the maximum number of nearby bookshops that a place can have while still being considered profitable. The second one, D, determines the maximum number of meters between two places that are considered near to each other.

The most important definitions contained in the original ontology used in this simulation are the following:

- An instance of class Distance should be related to two SpatialThing (the places between which the distance is computed) and to one integer (the meters between the two places):
  $\mathsf{Distance} \equiv = 2\, \mathsf{distBetween.SpatialThing} \sqcap = 1\, \mathsf{distMeasure}$

- To be considered near, two places should have a CloseDistance between them (a Distance which measures no more than D meters):
  $\mathsf{CloseDistance} = \mathsf{Distance} \sqcap \exists \mathsf{distMeasure}.(\leq, D)$

- A ProfitablePlace is a place that has no more than C bookshops nearby. In DL, it could be expressed as follows:
  $\mathsf{ProfitablePlace} \equiv \mathsf{SpatialThing} \sqcap$
  $\leq\ C\ \mathsf{hasDistance}.(\mathsf{CloseDistance} \sqcap \exists \mathsf{distBetween.bookshop})$

In order to learn the proper values to give to the adaptors C and D, a series of training axioms is fed to the system. More specifically, the city centres of the 30 largest city of the United Kingdom are classified as a ProfitablePlace and then additional information about each city is extracted using web data. This additional information describes places (such as bookshops) and the distances between them. To begin with, the location of a city centre is extracted from DBpedia [4]. A number of places around that city centre are then obtained from the Google Places API [14]. Some of these will be already classified Google as bookshops or as having a CloseDistance between them. The distances between them are then computed by the Google Distance Matrix API [13]. The resulting information was subject to noise; for example, several distances measuring more than ten kilometers were classified as CloseDistance. The sigmoid function used in the learning phase reduced the effect that values greatly differing from the mean have on the evolution of the adaptors.

Fig. 5 shows the evolution of adaptor D concurrent with the evolution of the adaptor C shown in Fig. 6. The ontology produced after the thirty iterations of this simulation defines a place to be profitable if there are no more than four bookshops within 1,360 meters distance.

A possible way to determine the correct value for the adaptor C is to consider the average number of bookshops near the city centres plus its standard deviation across the iterations (considering just the information contained in the training axioms). This value is found to be equal to 2.49. In a similar way the average measure of a CloseDistance plus the standard deviation is calculated as 1,325.

Assuming those values as the correct ones, the final value computed by the system for the adaptor D differs from the correct value by 4% of the standard deviation. The final value for adaptor C differs from the correct value by 162% of the standard deviation.

Fig. 5. Plot of the values of adaptor d across 30 iterations of the system

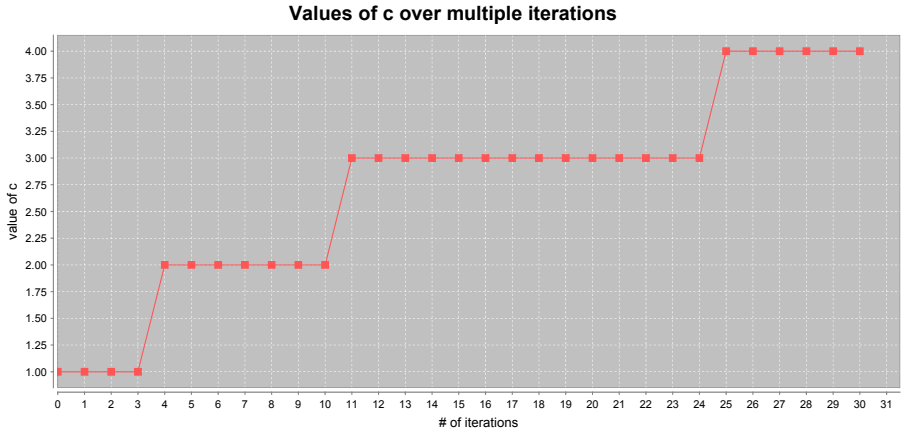

Fig. 6. Plot of the values of adaptor c across 30 iterations of the system

5 Related Work

A number of papers have addressed the problem of vagueness in ontologies with a common methodology of representing the indeterminacy of vague concepts as a static property, usually in terms of fuzzy logic and probabilistic techniques [25,30,5,23]. The approach we have presented here instead situates the problem of vagueness within the framework of ontology evolution and ontology alignment. That is, we focus on the dynamic properties of vague concepts, whereby there indeterminacy is bound up with their capacity to change and adapt to different contexts.

Unlike work in Ontology Learning (e.g., [36]), we are not attempting to construct ontologies from raw sources of information, such as unstructured text. Instead, our approach aims at computing ontological updates by aligning an existing ontology to another source of ontological information.

Several solutions have been proposed (e.g., [15]) to the problem of resolving inconsistencies in evolving ontologies, but none of them seem entirely satisfactory. One option is to develop reasoning methods that can cope with inconsistent axiom sets [26]. However, these methods are hard to automate and can be more time-consuming than traditional approaches. An alternative, proposed by [17], is to restrict updates to those that will preserve consistency. The disadvantage is that many kinds of ontology update will be disallowed, including the modification of vague concepts. Yet another strategy is to restore the consistency of an ontology (ontology repair) when an inconsitency arises. One possibility for automating the process of ontology repair is to remove some of the axioms that cause the inconsistency [16]. The axioms removed, however, might roll back the new changes that were introduced in the ontology or delete part of the ontology that should have been preserved. Our strategy for handling inconsistencies shares similarities with the ontology repair approach but differs from existing strategies in that no axioms are deleted as a result of the repair process.

6 Conclusion

The VAGO system presented here implements a novel approach for dealing with vagueness in formal ontologies: a vague concept receives a total interpretation (regarded as a supervaluation) but is inherently open to change through learning. More precisely, the meaning of a vague concept is dependent on a number of values, marked by adaptors, which can be automatically updated. These adaptors can be used to define cardinality restrictions and datatype range restrictions for OWL properties.

The definitions of the vague concepts of an ontology are automatically updated by validating the original ontology against a set of training axioms, thereby generating an updated ontology. Inconsistencies that arise from combining the training axioms with the the original ontology are interpreted as a misalignment between those two sources of ontological information. This misalignment can be reduced by modifying the values of the adaptors used in the original ontology.

If the axioms of another ontology are used as training axioms for the original ontology, then the update will result in an improved alignment between the two ontologies. The preliminary results obtained by the simulations suggest that this framework could be effectively used to update the definitions of vague concepts in order to evolve a single ontology or to improve the extension-based alignment between multiple ontologies.

References

1. Amoia, M., Gardent, C.: Adjective based inference. In: Proceedings of the Workshop KRAQ'06 on Knowledge and Reasoning for Language Processing. pp. 20–27. Association for Computational Linguistics (2006)
2. Baader, F., Hanschke, P.: A scheme for integrating concrete domains into concept languages. In: Proc. of the 12th Int. Joint Conf. on Artificial Intelligence (IJCAI91). pp. 452–457. Citeseer (1991)
3. Baader, F., Horrocks, I., Sattler, U.: Description logics. In: van Harmelen, F., Lifschitz, V., Porter, B. (eds.) Handbook of Knowledge Representation, chap. 3. Elsevier (2007)
4. Bizer, C., Lehmann, J., Kobilarov, G., Auer, S., Becker, C., Cyganiak, R., Hellmann, S.: Dbpedia - a crystallization point for the web of data. Web Semantics: Science, Services and Agents on the World Wide Web 7(3), 154 – 165 (2009)
5. Bobillo, F., Straccia, U.: fuzzyDL: An expressive fuzzy description logic reasoner. In: Fuzzy Systems, 2008. FUZZ-IEEE 2008.(IEEE World Congress on Computational Intelligence). IEEE International Conference on. pp. 923–930. IEEE (2008)
6. Brickley, D., Miller, L.: FOAF vocabulary specification 0.98, `http://xmlns.com/foaf/spec/`
7. Burato, E., Cristiani, M.: Learning as meaning negotiation: A model based on English auction. In: Håkansson, A. (ed.) KES-AMSTA 2009, pp. 60–69. No. 5559 in LNAI, Springer-Verlag (2009)
8. Choi, N., Song, I.Y., Han, H.: A survey on ontology mapping. SIGMOD Rec. 35, 34–41 (September 2006), `http://doi.acm.org/10.1145/1168092.1168097`
9. Drummond, N., Shearer, R.: The Open World Assumption. In: eSI Workshop: The Closed World of Databases meets the Open World of the Semantic Web (2006)
10. Euzenat, J., Shvaiko, P.: Ontology Matching. Springer, Berlin (2007)
11. Fine, K.: Vagueness, truth, and logic. Synthese 30, 265–300 (1975)
12. Flouris, G., Manakanatas, D., Kondylakis, H., Plexousakis, D., Antoniou, G.: Ontology change: Classification and survey. Knowl. Eng. Rev. 23, 117–152 (June 2008), `http://portal.acm.org/citation.cfm?id=1394822.1394823`
13. Google Inc.: The Google Distance Matrix API (Website accessed on: 12/08/2011), `http://code.google.com/apis/maps/documentation/distancematrix/`
14. Google Inc.: The Google Places API (Website accessed on: 12/08/2011), `http://code.google.com/apis/maps/documentation/places/`
15. Haase, P., van Harmelen, F., Huang, Z., Stuckenschmidt, H., Sure, Y.: A framework for handling inconsistency in changing ontologies. In: The Semantic Web – ISWC 2005, Lecture Notes in Computer Science, vol. 3729, pp. 353–367. Springer Berlin / Heidelberg (2005)
16. Haase, P., Stojanovic, L.: Consistent evolution of owl ontologies. In: Gómez-Pérez, A., Euzenat, J. (eds.) The Semantic Web: Research and Applications, Lecture Notes in Computer Science, vol. 3532, pp. 182–197. Springer Berlin / Heidelberg (2005)

17. Heflin, J., Hendler, J.: Dynamic ontologies on the web. In: Proceedings of the 17th National Conference on Artificial Intelligence and 12th Conference on Innovative Applications of Artificial Intelligence. pp. 443–449. AAAI Press (2000)
18. Horridge, M., Bechhofer, S.: The OWL API: A Java API for Working with OWL 2 Ontologies. In: OWLED 2009, 6th OWL Experienced and Directions Workshop (2009)
19. Hunter, J., McLaughlin, B.: JDOM, http://jdom.org/
20. Kamp, H.: Two theories of adjectives. In: Keenan, E. (ed.) Formal Semantics of Natural Languages. Cambridge University Press (1975)
21. Kennedy, C., McNally, L.: Scale structure, degree modification, and the semantics of gradable predicates. Language 81(2), 345–381 (2005)
22. Khattak, A., Pervez, Z., Lee, S., Lee, Y.: After effects of ontology evolution. In: Future Information Technology (FutureTech), 2010 5th International Conference on. pp. 1–6. IEEE (2010)
23. Koller, D., Levy, A., Pfeffer, A.: P-CLASSIC: A tractable probablistic description logic. In: Proceedings of the National Conference on Artificial Intelligence. pp. 390–397. Citeseer (1997)
24. Lehmann, F., Cohn, A.: The EGG/YOLK reliability hierarchy: Semantic data integration using sorts with prototypes. In: Proceedings of the third international conference on Information and knowledge management. p. 279. ACM (1994)
25. Lukasiewicz, T., Straccia, U.: Managing uncertainty and vagueness in description logics for the semantic web. Web Semantics: Science, Services and Agents on the World Wide Web 6(4), 291 – 308 (2008)
26. Ma, Y., Jin, B., Feng, Y.: Dynamic evolutions based on ontologies. Knowledge-Based Systems 20(1), 98–109 (2007)
27. Magka, D., Kazakov, Y., Horrocks, I.: Tractable extensions of the Description Logic $\mathcal{EL}$ with numerical datatypes. In: Automated Reasoning, Lecture Notes in Computer Science, vol. 6173, pp. 61–75 (2010)
28. Motik, B., Shearer, R., Horrocks, I.: Optimized reasoning in description logics using hypertableaux. In: Automated Deduction – CADE-21, Lecture Notes in Computer Science, vol. 4603, pp. 67–83 (2007)
29. Sirin, E., Parsia, B., Grau, B., Kalyanpur, A., Katz, Y.: Pellet: a practical OWL-DL reasoner. Web Semantics: science, services and agents on the World Wide Web 5(2), 51–53 (2007)
30. Stoilos, G., Stamou, G., Tzouvaras, V., Pan, J., Horrocks, I.: The Fuzzy Description Logic f-$\mathcal{SHIN}$. In: Proc. of the International Workshop on Uncertainty Reasoning for the Semantic Web. pp. 67–76. Citeseer (2005)
31. Van Deemter, K.: Not Exactly: In Praise of Vagueness. Oxford Univ Press (2010)
32. Van Rooij, R.: Vagueness and linguistics. In: Ronzitti, G. (ed.) Vagueness: A Guide. Springer Verlag (2011)
33. W3C OWL Working Group (ed.): OWL 2 Web Ontology Language Document Overview. W3C (October 2009), http://www.w3.org/TR/owl2-overview/
34. W3C Wiki: GeoOnion, http://www.w3.org/wiki/GeoOnion
35. Zablith, F., Sabou, M., d'Aquin, M., Motta, E.: Ontology evolution with Evolva. The Semantic Web: Research and Applications pp. 908–912 (2009)
36. Zhou, L.: Ontology learning: state of the art and open issues. Information Technology and Management 8, 241–252 (2007)